

CONSTRAINT ROUTE DISSEMINATION USING
DISTRIBUTED ROUTE EXCHANGERS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 This is the first application filed for the present invention.

MICROFICHE APPENDIX

Not Applicable.

TECHNICAL FIELD

10 The present invention relates in general to routing control in data networks, and in particular to constraint route dissemination in a data network that uses an Open Shortest Path First (OSPF) routing protocol.

BACKGROUND OF THE INVENTION

15 Open Shortest Path First (OSPF) is a routing protocol developed for Internet Protocol (IP) networks based on the shortest path first or link-state algorithm.

Routers use link-state algorithms to send routing information to all nodes in an IP network by calculating the shortest path to each node based on a topography of the IP network constructed by each node. Each router maintains a routing table in which the router stores routes to particular network destinations. Each router sends the state of its own links to other routers in the IP network.

25 The advantage of shortest path first algorithms, compared to distance vector algorithms, is that they result in smaller more frequent updates everywhere in the IP

09704291-10100
00707T-16240260

network. They also converge quickly, so that problems such as routing loops and infinite counts, in which routers continuously increment the hop count to a particular destination node, are avoided.

5 A disadvantage of OSPF algorithms is that they require a great deal of processing power and memory in the OSPF routers. However, it is generally believed that the advantages outweigh the disadvantages, and OSPF Version 2, which is defined by RFC 1583, appears to be rapidly
10 replacing Reservation Initiation Protocol (RIP) on IP networks.

Current OSPF link-state routing protocols flood Link-State Advertisements (LSAs) to all OSPF routers. This provides OSPF routers with routing information for routing
15 payload packets on a hop-by-hop basis. However, for the purposes of establishing paths for Traffic Engineering (TE), TE routing information is only required at a point where the path setup is triggered. Consequently, it is not necessary for all routers in a routing area to flood,
20 process and store a large number of TE-LSAs.

If each OSPF router in an IP network of a number (n) of OSPF routers originates a number (L) of LSAs and floods the LSAs on a number (p) of ports, each router in the network potentially receives up to $n \cdot L \cdot p$ LSAs. The
25 total number of LSAs processed by all OSPF routers in the network is therefore potentially $n \cdot n \cdot L \cdot p$ LSAs. This places an unnecessarily large computation load on the OSPF routers, and forces the provision of high speed processors and a large amount of memory in the routers. It also

09704291 110100

appropriates network bandwidth for LSA exchange that could be more profitably used for payload packet traffic.

FIG. 1 is a schematic diagram of an Internet Protocol (IP) network 10 that uses an Open Shortest Path First (OSPF) routing protocol in accordance with the prior art. The network also includes a plurality of routers 12-22, labeled as routers A-G. The structure and functions of the network 10 are well known in the art.

10 SUMMARY OF THE INVENTION

The invention provides a method and apparatus for establishing explicit edge-to-edge paths in an Internet Protocol (IP), Multi-Protocol Label Switched (MPLS) or Optical network that uses an open shortest path first (OSPF), or similar routing protocol, for computing explicit constrained paths.

The invention adds a new type of router to the network. The router is a traffic engineering route exchange router (TE-X) that maintains a traffic engineering link-state database (TE-LSDB) using link state information sent by associated OSPF routers in traffic engineering link state advertisement (TE-LSA) messages. The association between a TE-X and other OSPF routers is preferably based on proximity (least cost). The OSPF router only sends TE-LSAs to its associated TE-X via a fast path. The TE-LSAs are not flooded to all routers in the network.

In accordance with the method, traffic engineering link state advertisement (TE-LSA) messages are sent from

09704291.110100

OSPF routers in the network to a nearest traffic engineering route exchange router (TE-X) in the network, to permit the TE-Xs in the network to maintain traffic engineering link-state databases (TE-LSDB). The nearest
5 TE-X is queried to obtain an explicit edge-to-edge path satisfying specified traffic engineering (TE) parameters supplied by the querying OSPF router.

It should be understood that the concept of the TE-X and the mechanisms for discovering and distributing
10 routes in accordance with the invention are applicable to other routing protocols besides OSPF. For example, Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-Mode Network
15 Service (ISO 8473), International Standard ISO 10589, 1992.

The invention therefore provides an improved data network in which traffic engineering message volume is reduced and the computing power and memory required for most routers in the network is dramatically reduced.
20 Network resources are therefore conserved and network provisioning is simplified.

BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages of the present invention will become apparent from the following detailed
25 description, taken in combination with the appended drawings, in which:

FIG. 1 is a schematic diagram of an OSPF network in accordance with the prior art;

09704291 "110100

FIG. 2 is a schematic diagram of an OSPF network in accordance with the invention;

FIG. 3 is a schematic diagram of the OSPF network shown in FIG. 2, illustrating traffic engineering message exchange between the OSPF routers in the network and traffic engineering route exchange routers in accordance with the invention; and

FIG. 4 is a state and diagram illustrating the states of the traffic engineering route exchange routers in accordance with the invention.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

FIG. 2 is a schematic diagram of an Internet Protocol (IP) network 30 in accordance with the invention that uses the OSPF routing protocol. The network 30 includes a plurality of interconnected routers 32-40 and three traffic engineering route exchange routers (TE-X) 42-46, for the sake of example. As will be explained below in more detail, the TE-Xs 42-46 receive Traffic Engineering Link-State Advertisement messages (TE-LSAs) from the routers 32-40 and use the LSAs to construct a traffic-engineering database (TE-database) 48-52. The TE-databases 48-52 are used to store explicit routes between edge routers (32 and 40, for example) in the IP network 30.

FIG. 3 schematically illustrates the TE-LSA message exchanges in the IP network 30 in accordance with the invention. TE-LSAs sent from the OSPF routers 32-42

[illegible]

are shown as solid lines 60, whereas TE-LSAs sent from TE-Xs to other TE-Xs to synchronize the TE-databases 48-52 are shown as dashed lines 70. As will be explained below in more detail, TE-LSAs are sent directly from OSPF routers on a fast path to a closest TE-X 42-46. Each OSPF router 32-40 in the network 30 originates a number (L) of TE-LSAs and sends the L TE-LSAs directly to a TE-X 42-46, only the TE-Xs 42-46 collectively receive the $n \cdot L$ LSAs, where n is the number of OSPF routers in the network 30.

Note that the L TE-LSAs are not flooded on all ports of the respective routers 32-42. The total number of TE-LSAs processed in the network 30 is therefore $x \cdot n \cdot L$, where x is the number of TE-Xs 42-46 in the network 30. The invention therefore reduces the total number of TE link states flooded, processed and stored in a network. In other words, the invention reduces the total amount of network and router resources required for processing TE-LSAs, from $O(n^2)$ to $O(n)$, in a network 30 of n nodes.

To setup a constraint path in the network 30, an edge router (32, for example) queries the nearest TE-X (42 in this example) for a route satisfying specified constraints. The TE-X 42 returns an explicit set of routes (an Explicit Route Object (ERO), as specified by the Multi-Protocol Label Switched Path (MPLS) protocol). The exact semantics of the query and response message are not a part of this invention.

An edge router discovers TE-Xs in the network 30 via normal Router LSAs, as will be explained below in more detail. TE-Xs in a routing area elect a Designated TE-X and peers with it. Each Area Border Router (ABR), which must also be a TE-X, peers with other TE-Xs in other areas,

09704291.110100

in the same way. The routers in the network are assumed to be connected (e.g. via a control channel) and can be reached via shortest path routes computed using normal LSAs flooded by OSPF.

5 As explained above, the TE-LSAs are not flooded like normal LSAs used to calculate the "shortest path". Consequently, a router 32-40 in the network 30 only sends out one copy of its TE-LSAs and the TE-LSAs are forwarded directly (on a fast path) to the nearest TE-X, without
10 incurring processing delay (in addition to slow path delay) on every hop, which occurs when normal LSAs are flooded. Similarly, a TE-X sends TE-LSAs directly to other TE-Xs via the mechanism of peering with other TE-Xs. As a result, the time required for TE link states to converge within the
15 routing domain is reduced, and routers not involved in initiating the establishment of constraint paths are not burdened with processing and storing (TE/constraint) link states (in addition to normal link states) that they do not need. Only "edge" routers need to have access to the
20 explicit constraint routes.

It should be noted that only routers (OSPF-TE) that support constraint path setup need to be provisioned with the functionality of a TE-X 42-46 in accordance with the invention. It is possible to upgrade existing OSPF
25 label switching routers (LSRs) with OSPF-TE support since there is very little processing and storage impact on the LSRs (see "draft-ietf-mpls-arch-07"). The heavier processing and storage requirement is at the TE-X, which must process and store a large number of TE-LSAs and
30 compute TE routes.

09704291-110100

In accordance with the invention, the TE-LSAs are sent directly to the "nearest" TE-X. TE-Xs are discovered through normal Router Link State Advertisement (LSA). A TE-X exchanges TE-LSAs with other peer TE-Xs. Edge routers
5 query a TE-X for a constraint or TE route to a destination.

To be able to discover TE resources (e.g., links, bandwidth) to another router C (router 34 in FIG. 3), a router A (router 32 in FIG. 3) must already have connectivity to router C, either via an existing link or a
10 control channel setup for control messages (for e.g. MPLS signaling, or OSPF control messages) between C and A.

If C is reachable (but not directly), it is also possible for A to learn (through the TE-X) about links that could potentially be used to reach C directly. It is
15 likewise possible for A to advertise (again via a TE-X) the potential links that connect it directly, or in 'one hop', to other networks that may include another router, and vice-versa. This is useful in applications like virtual private network (VPN), where a router may have connectivity
20 to another router, but not a direct link, and may setup "on-demand" circuits to the other router.

Currently, this is not feasible (every TE link must have a corresponding control channel) because OSPF routers have to flood and exchange TE-LSAs with their
25 neighbors. As explained above, in accordance with the invention OSPF routers send TE-LSAs to a nearest TE-X only. Hence, it is not mandatory to have a control channel for every set of TE links to enable TE routing information to be disseminated.

Changes to OSPF Routers Originating TE-LSAs

In accordance with the invention, OSPF routers are routers 34-38 that are not TE-X 42-46 routers. OSPF routers that support TE-LSA distribution are referred to as
5 OSPF-TE routers. The normal Link-State Database is referred to as a LSDB and the Engineering Link-State Database 48-52 (FIG. 3) are referred to as TE-LSDBs.

OSPF routers that originate TE-LSAs send Link-State Update (LSU) messages directly via unicast to the
10 nearest TE-X instead of flooding them from each of their interfaces. TE-LSAs are only received by TE-Xs. OSPF-TE routers do not receive TE-LSAs originated by other OSPF routers, nor do they send TE-LSAs to other OSPF routers (unless they are TE-Xs) or exchange TE-LSAs with other OSPF
15 routers.

The behavior of OSPF routers with respect to normal LSAs is not changed, with the exception of the processing of Router-LSAs. When a new or newer Router-LSA instance is received (as described in [OSPF]), and the
20 Option Field of the outer-LSA indicates that it is a TE-X, as will be explained below, the OSPF-TE router uses the instance to maintain a list of TE-Xs in the routing area. It also uses the Router-LSA instance to compute an entry in a normal OSPF routing table for the Advertising Router,
25 which is the Router ID for a Router-LSA. The entry in the OSPF routing table is similar to an Autonomous System Border Router (ASBR) route entry. The entry in the routing table is in addition to the normal installation procedure of the Router-LSA in the normal LSDB, which is well known
30 in the art. If the Option Field does not have the TE-X indication set, the Router-LSA is processed and installed

09704291.110100

as described in RFC2328. The Router ID of a TE-X should be a unique IP address identifying the TE-X router in an Autonomous System, and is preferably not an interface address, but rather an address not assigned to any other
5 interface. It may, however, be an address assigned to a "loopback" or dummy interface.

When an OSPF router is initialized, it establishes
adjacencies with its neighbor(s). Once the adjacencies are
established (Section 7 of OSPF), an OSPF-TE router
10 determines the nearest TE-Xs from the normal routing table
it has computed (See Section "Computing TE-X Routing Entry"
below) for the Link-State Database (LSDB). The OSPF-TE
goes through its list of TE-Xs, and preferably caches up to
two TE-Xs (the primary TE-X and a backup TE-X), as
15 described below in the Section "Determining Nearest TE-Xs".
TE-X capability is advertised in Router-LSA instances, as
described below in the Section "Advertising TE-Xs".

Originating TE-LSAs

20 Prior to sending a TE-LSA, the OSPF-TE router
should send a TE Database summary list and wait for
acknowledgement from the TE-X. The details of this
procedure are described in below in the section "Sending
Link State Update TE-LSAs".

25 The OSPF-TE router then proceeds to originate
TE-LSAs describing its connectivity or links and the
constraints associated with the connectivity, and sends
these TE-LSAs to the nearest TE-X. The TE-LSAs are defined
in [OSPF-TE-EXT], and have the following format:

```

5      +-----+
      | Standard LSA Header |
      +-----+
      | Type, Length, Value |
      +-----+

```

There are 2 types currently defined:

- 10 1) Router Type, Length, Value (TLV) - describes the address of the router that is always reachable;
- 2) Link TLV - describes a single link.

The following sub-TLVs of the Link TLV - Link type, Link ID, local/remote interface IP address describes the link and the maximum bandwidth, maximum reservable bandwidth, resource class, describes the constraints of that link. Additional sub-TLVs are defined in *ompls* and these TE-LSAs can be distributed using the mechanisms in accordance with the invention.

20 If the TE-X does not acknowledge a TE-LSA sent, the OSPF-TE router must attempt to retransmit the TE-LSA to the backup TE-X.

VPN Sub-TLV

25 The purpose of the VPN sub-TLV is to indicate to other nodes in the network the VPN Identifier (VPN ID) [RFC2685], that is associated with a router.

The format of the VPN sub-TLV is as shown below in

30 Table 1:

09704291.110100

TABLE 1

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
5 |                                     | Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     No. VPN Attachments |T|
+-----+-----+-----+-----+-----+-----+-----+-----+
| Reserved |
+-----+-----+-----+-----+-----+-----+-----+-----+
10 |                                     VPN ID
+-----+-----+-----+-----+-----+-----+-----+-----+
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The VPN sub-TLV is optional. The VPN ID is as defined in [RFC 2685], is 7 octets in length. The "T" field indicates whether a router is connected to a transit VPN node. A transit VPN node permits other VPN nodes to connect to it to reach other VPN nodes. The "No. VPN Attachments" (number of VPN attachment allowed) indicates the maximum number of VPN nodes that can be connected to a transit VPN node, in order to reach other VPN nodes having the same VPN ID.

Replicating or Multicast Capable Sub-TLV

25 The purpose of the Replicating capable sub-TLV is
to indicate to other nodes that a router is capable of
replicating data to more than one end point.

The format of the Replicating Capable sub-TLV is shown below in Table 2:

TABLE 2

```

30      0      1      2      3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
      |                                     |                                     Length                                     |
      +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
35      |                                     No. of Receivers                                     |
      +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

The "No. of Receivers" field indicates the number of receivers permitted to be attached to a Replicating Capable router.

Discovering TE-Xs

5 OSPF-TE routers discover TE-Xs via normal Router-LSA flooding, as described below with reference to "Advertising the TE-X". When an OSPF-TE router receives a Router-LSA with the TE-X indication bit on, it stores an IP address of the TE-X in a TE-X data structure. In addition,
10 once a Dijkstra computation of routes to all destinations in the network is performed, the OSPF-TE also stores a cost of a shortest path route to each TE-X in a designated area of the TE-X data structure.

15 Determining Nearest TE-Xs

 The "nearest" TE-X to an OSPF-TE is the TE-X with a least total link cost from the OSPF-TE, which is computed as described above. If an OSPF-TE router is a TE-X, then the nearest TE-X is the router itself. Once an OSPF-TE has
20 received the routing information in the domain, it goes through its list of TE exchanges, starting from the nearest TE-X, to find out which of the TE-Xs will agree to serve the OSPF-TE. The route to a TE-X is determined from the routing table calculation described in Section 16.1 of
25 [OSPF]. In accordance with the invention, the OSPF-TE sends a probe message to each of the TE-Xs in the list, until two TE-Xs agree to serve the OSPF-TE by acknowledging the probe message. The nearest (least cost) TE-X acknowledging the probe message is designated as the
30 primary TE-X, and the next nearest TE-X acknowledging the

09704291.110100

probe is designated as the backup TE-X. If two TE-Xs have the same route cost, the one with the smallest difference in address value compared to the Router ID of the OSPF-TE, is chosen as the primary TE-X and the other as the backup
5 TE-X.

If only one TE-X responds to the probe messages, then the OSPF-TE is served by only one TE-X, and the OSPF-TE preferably generates an alert message to a network router management system. In an operational network, it
10 should normally be sufficient to have three distributed TE-Xs. Consequently, if one TE-X fails, a primary and secondary TE-X are always available to serve all the OSPF-TEs, assuming the network is not partitioned.

15 **Sending Link State Update TE-LSAs**

OSPF routers send TE-LSAs directly to the nearest TE-X. If an OSPF router does not receive a Link State Acknowledge for a TE-LSA sent to a particular TE-X after a timeout period, it sends the TE-LSA to a next nearest TE-X
20 and re-invokes the "Determining nearest TE-X" mechanism to discover current serving TE-Xs.

Initially, a Database Summary List is sent to the TE-X, and the sending OSPF-TE expects to receive Link State Requests related to the TE-LSAs that the TE-X wants to
25 update in its database. The OSPF-TE sends the TE-LSAs to the TE-X as requested in a Link State Request message. The difference between this procedure and OSPF database synchronization is that the OSPF-TE does not request TE-LSA from the TE-X.

30 An OSPF router does not initiate adjacency establishment nor maintain adjacency or exchange hellos

09704291.110100

with a TE-X. An OSPF router is able to discover the serving TE-X via Router-LSAs, as explained above with reference to "Determining Nearest TE-Xs".

5 TE Exchange

TE-Xs must participate in normal OSPF route distribution, although they may not necessarily be participating in path setup or be able to "label switch", for instance a TE-X may be a leaf of the IP network 30, setup to function solely as a traffic engineering route exchange router (TE-X).

Advertising the TE-X

A TE-X must advertise its capability. Advertising of the TE-X is preferably accomplished using the Router LSA "Options" field as shown below. A TE-bit in the Options field is defined for this purpose. It should be noted that a TE-X also indicates whether it is an area border router, or an AS boundary router, by setting the appropriate bits (bit B and bit E, respectively) in its router-LSAs.]

Options Field in LSA Header is shown in Table 3:

09704291.110100

Table 3

The Options Field

```

+-----+
| TE | O | DC | EA | N/P | MC | E | * |
+-----+

```

5

*Reserved bit

As seen in Table 3, the options field contains 8 bits, 6 of which have to date been reserved for the purposes described below. In accordance with the invention, one of the two reserved bits is used as a TE-bit for advertising a TE-X. The bits in the Options field are respectively used as follows:

E-bit

15 Specifies the way AS-external-LSAs are flooded, as described in Sections 3.6, 9.5, 10.8 and 12.1.2 of [OSPF].

MC-bit

20 Specifies whether IP multicast datagrams are forwarded according to the specifications in [MOSPF].

N/P-bit

Specifies the handling of Type-7 LSAs, as specified in [NSSA].

DC-bit

25 Specifies the router's handling of demand circuits, as specified in [DEMD].

EA-bit

30 Specifies the router's willingness to receive and forward External-Attributes-LSAs, as specified in [EAL].

09704291.110100

O-bit

Specifies the router's willingness to receive and forward Opaque-LSAs as specified in [rfc2370].

TE-bit

- 5 Specifies a Traffic Engineering Route Exchange Router in accordance with the invention for receiving and distributing constraint routes, as well as computing and providing constraint paths or explicit routes when queried.

10

TE Exchange Peering

- When a TE-X is initialized, the TE-X behaves like a normal OSPF-TE, creating adjacencies with its neighbor(s) [designate routers (DRs) and backup designated routers (BDRs) on broadcast and NBMA networks) to exchange normal (non TE) routing information. Once a TE-X has established adjacencies and downloaded the domain link state database (as defined in the RFC2328), the TE-X begins to establish adjacencies or peering with other TE-Xs in the area, that it learned from normal (non-TE) Router-LSAs (and stored in the list of TE-Xs). The peering with other TE-Xs is established similarly to the creation of adjacencies with OSPF neighbors. The TE-X sends (unicast) to each other TE-X in the routing area Hello or Keep-Alive messages (to differentiate from OSPF Hello messages). Once the TE-X DR of the routing area is discovered via the Keep-Alive messages, the TE-X attempts to establish an adjacency with the TE-X DR to exchange TE routing information. Once bi-directional communication has been established with the TE-X DR (i.e. 2-Way Received neighbor state), the TE-X proceeds to exchange the TE-LSAs, in the same way as

09704291-110100

A designated router TE-X, the TE-X DR, and the backup designated router TE-X, the TE-X BDR, in a routing area are preferably elected using the same procedures as are used to elect an OSPF DR and BDR. The TE-X DR is responsible for peering with other TE-Xs in a routing area. All TE-Xs maintain a consistent TE-LSDB of the routing area by synchronizing and exchanging TE-LSAs. It should also be noted that a TE-X does not peer with other (non TE-X) OSPF-TEs.

15 Once an OSPF-TE router has setup a path (the
router may be functioning as a Label Switched Router, LSR),
it should update the designated TE-X of its currently
available resources. To accomplish this, the OSPF-TE
originates new TE-LSAs reflecting its available resources.
20 On receipt of the new TE-LSAs, the TE-X distributes the new
TE-LSAs to each of its peer TE-Xs.

There may be instances during setup of an explicit path (LSP1) in which a TE-X, using its existing TE-LSDB, computes a path that traverses links that do not have enough free resources, because a conflicting path is being established. If this happens, the path setup algorithm may choose to use crank-back, and attempt to obtain a diverse explicit route to the destination from the point where the path setup was cranked back. The router may obtain the diverse explicit route from the TE-X. The choice of heuristics in path setup is beyond the scope of this

invention, however, an example of path setup is described to illustrate how routers may obtain an explicit route from a TE-X, under certain circumstances.

5 The necessity of having absolute up-to-date information should be carefully examined. Nevertheless, the timely distribution of available resources can be improved using the following mechanism. When an edge router 32 (FIG. 3) requests a 2.5 Gbs constraint shortest path to and another edge router 40, for example, the
10 TE-X 42 may originate the delta of resources reserved (RR), RR TE-LSAs (in this case the bandwidth reserved along the C-SPF) and distributes the RR TE-LSAs to other peering TE-Xs 44 and 46 via messaging channels 70. This prevents the other TE-Xs 44,46 from assigning routes along the same
15 path assigned by TE-X 42, if inadequate resources exist along a link in the path after the path is setup along the constrained shortest path first (C-SPF) route computed by TE-X 42. RR TE-LSAs reduce the occurrence of "resource allocation collision", by leveraging each TE-X's a priori
20 knowledge of resources allocated, if the knowledge is rapidly distributed to other TE-Xs.

As described above, all TE-Xs must have consistent TE-LSDBs, but OSPF-TEs do not maintain TE-LSDBs. RR TE-LSAs are only originated by TE-Xs and reflect the delta
25 of resources used. However, RR TE-LSAs do not change the TE topology database. When the connectivity associated with constraints or traffic engineering of an OSPF-TE router changes (e.g. due to, interface failure or link failure), the OSPF-TE must originate new TE-LSAs reflecting
30 the connectivity and currently available resources. A TE-X must flush any corresponding RR TE-LSA from its TE-LSDB

09704291.110100

upon receiving a more recent TE-LSA from an OSPF-TE. Other peering TE-Xs perform the same action when they receive a copy of the more recent TE-LSA. This ensures that each TE-LSDB in the routing area is consistently up to date. If a path setup fails or returns an error condition, or if a path is torn down, the router requesting the explicit route from the TE-X, should send a "release" explicit route message, to enable the TE-X to originate new "RR" TE-LSAs to flush the older RTE-LSA from the TE-LSDB of peer TE-Xs.

10 If the resources allocated by routers during path setup are different from the resources requested in the path setup, OSPF-TE routers should also originate new TE-LSAs to reflect the change.

15 **Loose Source Routing**

A Label Switching Router (LSR) may obtain an explicit route from a TE-X if it is provided with a Loose Source Route in a path setup-signaling message.

20 **Area Border Routers (ABRs)**

Traffic engineering area border routers (TE-ABRs) send out network TE-LSA summaries. The definition of the network summary TE-LSA is beyond the scope of this invention. However, each TE-ABR in a routing area should be a TE-X in each routing area to which it is connected, to permit the ABR to exchange summary TE-LSAs with peer TE-Xs in other routing areas. This has the advantage of permitting a TE-X that is an ABR to provide information respecting paths across another area, and available resources associated with the paths. Otherwise, edge routers are not aware of resources available on routers in

09704291-110100

other areas. If it is not feasible to have every ABR function as a TE-X, an alternative is to have a domain router TE-X (DR TE-X) in a routing area summarize the network LSAs in that area and distribute the summary to all the ABRs in the routing area. The ABRs should distribute those summary TE-LSAs to the DR TE-X in other routing area(s) to which each ABR is connected. All ABRs in a routing area send network summary TE-LSAs received from other routing areas to the DR TE-X in their area. The DR TE-X in that area should in turn distribute the network summary TE-LSAs from other routing areas to all TE-Xs in the routing area.

TE-X List

Each OSPF-TE maintains a TE-X list for each routing area. An alternate option is to use OSPF Opaque LSAs to distribute router capability to scope the flooding of TE-X capability to a routing area, subnet or autonomous system (AS). A disadvantage is that extra OSPF Opaque LSAs need to be flooded. However, by setting a bit in the Options Field of the Router-LSAs to indicate the router capability, the number of control messages is not increased. It also possible to have different TE-Xs process different Resource Colors, but this is more involved, and more TE-X adjacencies are required. The concept can be expanded to allow different classes of constraints, for example, links for low delay traffic need only be sent to TE-Xs that serve low delay routes.

FIG. 4 is a state diagram that schematically illustrates the various states of a TE-X 42-46 in accordance with the invention. In a principal state 90,

09704291-110100

the TE-X 42 is waiting for a message. When a message is received, the message is deciphered in a message deciphering state 92. Different message types require different actions. Certain message types require only an acknowledgment. Acknowledgments are processed in state 94 and the TE-X returns to the state 90 to wait for messages, if the message processed in state 92 requires no further action. If the message is a TE-LSA, the TE-X analyzes the TE-LSA in state 96 to determine whether the TE-LSA is valid. If the TE-LSA is determined to be valid, the TE-X forwards in the TE-LSA to its peer TE-Xs. Otherwise, the TE-X enters state 104, in which it reports the invalid TE-LSA to a network management authority. The TE-X then returns to the waiting for message state 90.

The TE-X also receives requests for explicit paths from OSPF-TE routers, as explained above. If the message received is a request for an explicit path, the TE-X enters state 100 in which it computes an explicit path using its TE-LSDB, as explained above. The TE-X then enters state 102, in which it returns a path response message to the requesting OSPF-TE. The TE-X then returns to state 90 in which it waits for messages.

The embodiment(s) of the invention described above is(are) intended to be exemplary only. The scope of the invention is therefore intended to be limited solely by the scope of the appended claims.

09704291-110100